Is this word borrowed? An automatic approach to quantify the likeliness of borrowing in social media

[1]Jasabanta Patro, [2]Bidisha Samanta, [3]Saurabh Singh,
[4]Prithwish Mukherjee, [5]Monojit Choudhury, [6]Animesh Mukherjee
[1,2,3,4,6] Indian Institute of Technology Kharagpur, India – 721302 , [5]Microsoft Research India, Bangalore – 500033
[1]jasabantapatro, [2]bidisha, [3]saurabhsingh}@iitkgp.ac.in,
[4]pritspido@gmail.com, [5]monojitc@microsoft.com, [6]animeshm@cse.iitkgp.ernet.in

arXiv:1703.05122v1 [cs.CL] 15 Mar 2017

Abstract

Code-mixing or *code-switching* refer to the phenomenon of effortless and natural switching between two or more languages in a single conversation, sometimes even in a single utterance, by multilingual speakers. However, use a foreign word in a language does not necessarily mean that the speaker is code-switching, because often languages *borrow* lexical items from other languages. If a word is borrowed, it becomes a part of the lexicon of a language; whereas, during code-switching the speaker is aware that the conversation involves multiple languages and often the switching is intentional. Identifying whether a non-native word used by a bilingual speaker is due to borrowing or code-switching is not only of fundamental importance to theories of multilingualism, but it is also an essential prerequisite towards development of language and speech technologies for multilingual communities.

In this paper, we present for the first time, a series of computational methods to identify the likeliness of a word being borrowed or code-mixed, based on the signals from social media. In particular, we use tweets from English-Hindi bilinguals from India to predict word borrowing. We first propose a method to sample a set of candidate words from the social media data using a context based clustering approach. Next, we propose three novel and similar metrics based on the usage of these words by the users in different tweets; we then apply these metrics to score and rank the candidate words indicating their likeliness of being borrowed. We compare these rankings with a ground truth ranking constructed through a human judgement experiment. The Spearman's rank correlation between the two rankings (~ 0.62 for all the three metric variants) is more than double the value (0.26) of the most competitive existing baseline reported in the literature. Some other striking observations are – (i) the correlation is higher for the ground truth data elicited from the younger participants (age < 30) than that from the older participants; since language change is brought about by the younger generation, this possibly indicates that social media is able to provide very early signals of borrowing, and (ii) those participants who use mixed-language for tweeting the least, provide the best signals of borrowing.

Introduction

In multilingual societies, where two or more languages coexist and are used regularly by the users, several language-interaction phenomena are observed. *Code-switching* or *code-mixing* is one such common phenomenon where speakers spontaneously switch between multiple languages in a single conversation, sometime even within a single sentence or phrase (Auer 1984). A related but linguistically and cognitively distinct phenomenon is *lexical borrowing* (or simply, *borrowing*), where a word or phrase from a foreign language is used as a part of the native vocabulary of a language. Examples of word borrowing are widespread; for instance, in Dutch the English word "sale" is now used more frequently than the Dutch equivalent "uitverkoop". Some English words like "shop" are even inflected in Dutch as "shoppen" and heavily used. Similar examples of borrowing of words from English can be found across many other languages like Spanish, German, French, Chinese etc. The opposite is also true, i.e., words like "tortilla", "tequila", "ramedan", "couscous" and "tandoori" are borrowed in English from other languages.

While it is difficult in general to ascertain whether a foreign word or phrase used in an utterance is borrowed or just an instance of code-mixing (Bali et al. 2014), one tell tale sign is that only proficient multilinguals can code-mix, while even monolingual speakers can use borrowed words because, by definition, these are part of the vocabulary of a language. In other words, just because an English speaker understands and uses the word "tortilla" does not imply that she can speak or understand Spanish. In this work, we develop a novel method to identify borrowed words, or more formally the likeliness that a word is in the process of getting borrowed.

A borrowed word from a foreign language, initially appears frequently in speech, then gradually in print media like newspaper and finally it loses its origin's identity and is used in the native language resulting in an inclusion in the dictionary of the native language (Myers-Scotton 2002; Thomason 2003). However, early-stage automatic identification of whether a word is likely to be borrowed is known to be a hard problem. The main hurdles are (i) the lack of well-defined linguistic signals of borrowing, primarily because it is a socio-linguistic phenomenon closely related to acceptability and frequency, (ii) borrowing is a dynamic process; new borrowed words enter the lexicon of a language as old words, both native and borrowed, might slowly fade away from usage, and (iii) it is a population level phenomenon that

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Figure 1: Some example code-mixed tweets from English-Hindi bilinguals. Hindi words are in italics.

Huge traffic restrictions for PM's visit to #blast site mean deserted roads in #Hyderabad. ''*Itna sanaata kyon hai bhai?*''
Translation: Huge traffic restrictions for Prime Minister's visit to the blast site mean deserted roads in Hyderabad. "Why is there so much silence, bro?"

MMS will go to #HyderabadBlast site to take *jayeja* of area & say *Hazaaron Jawabon Se Acchi Hai Meri Khamoshi* #ThikHai
Translation: MMS (name of a politician) will go to #HyderabadBlast site to take a survey of the area and say "My silence is better than a thousand answers." #ThikHai

necessitates data from a large portion of the population unlike standard natural language corpora that typically comes from a very small set of authors.

The above reasons motivate us to resort to the social media (in particular, Twitter), where a large population of bilingual/multilingual speakers are known to often tweet in code-mixed colloquial languages (Carter, Weerkamp, and Tsagkias 2013; Solorio et al. 2014; Vyas et al. 2014). Fig. 1 shows some typical tweets from Hindi-English bilinguals from India.

The central hypothesis of this study is as follows: Since, language use over social media is informal, has speech-like characteristics (Pérez-Sabater 2012) and involves a large population of speakers from a wide range of socio-linguistic communities, it should be possible to extract early signals of likeliness of borrowing of a word from the language usage patterns from social media data. In this study, we analyze the tweets of the English-Hindi bilinguals from India for predicting the likeliness of borrowing. Note that our approach and the metrics that we propose can be easily generalised for any other language pair. Hindi is assumed to be the native language while English is the foreign language from where a candidate word could be borrowed. We present some typical examples of such bilingual tweets in figure 1.

The main stages of our research are as follows:

- *Unsupervised method for selecting candidate words from a large corpora*: We propose a context based clustering approach to appropriately sample a list of candidate English words from the bilingual tweets to investigate the borrowing phenomenon.
- *Ground truth generation*: We launch an extensive survey among 58 human judges of various age groups and various educational backgrounds to collect responses indicating if each of the candidate English word is likely borrowed. Further, we aggregate the responses to rank the candidate words based on the likeliness of borrowing.
- *Metric to quantify the likeliness of borrowing from social media signals*: We define three novel and closely similar metrics that serve as social signals indicating the likeliness of borrowing. All the three metrics attempt, in some form, to estimate the extent to which users use a candidate English word in an otherwise (predominantly) Hindi tweet. The higher this extent is for a candidate word, the higher should be the likeliness of borrowing.
- *Experiments*: We compare the likeliness of borrowing as predicted by our model and a baseline model with that of ground-truth obtained from human judges.

Finally, our key results are outlined below:

- The Spearman's rank correlation between the ground-truth ranking and the ranking based on our metric is ~ 0.62 for all the three variants. Remarkably, this value is more than double the value (0.26) if we use the most competitive baseline (Bali et al. 2014) reported in the literature for ranking the words.
- In case the candidate words are surely instances of borrowing (i.e., are at the top of the ground truth ranking), our metrics does as good as the baseline. However as one moves down the rank list, our metric overwhelmingly outperforms the baseline.
- Interestingly, the responses of the judges in the age group below 30 seem to correspond even better with our metrics. Since language change is brought about mostly by the younger population, this might possibly mean that our metrics are able to capture the early signals of borrowing.
- Those users that mix languages the least in their tweets present the best signals of borrowing in case they do mix the languages (correlation of our metrics estimated from the tweets of these users with that of the ground truth is ~ 0.65).

Note that apart from the fact that this is a hard but very interesting socio-linguitic problem (which was our primary motivation to choose this problem), the ability to automatically distinguish cases of code-borrowing from code-mixing can have a strong impact on engineering applications like multilingual information retrieval and natural language processing tasks. For instance, if a query has more than one language, say L_1 & L_2, and if all the words of L_2 are borrowed in L_1 and not vice versa, then we can infer that it is possibly an L_1 query, thus giving priority to only L_1 documents. However, if L_1 and L_2 do not have mutually borrowed words, then it is a truly multilingual query and so, we should retrieve L_1, L_2 and L_1-L_2 mixed documents for that query. Furthermore, several studies have brought up into notice the prevalence of multilinguality, and more specifically use of code-mixing in user-generated content on social media (Bali et al. 2014; Barman et al. 2014; Solorio et al. 2014). Processing of such content, for example sentiment and opinion detection of code-mixed tweets (Rudra et al. 2016), requires language detection, and more specifically language-switch detection because language switch often signals a change in opinion or sentiment. However, as argued in the case of IR, borrowed words do not necessarily indicate language switch, even though they belong to a different language.

Related Work

Early literature: In linguistics, *code-mixing* and *word-borrowing* are often studied under the broader scope of language evolution and change. Both of these phenomena are observed when the speakers of two or more languages come in close contact with one another. Linguists have for a long time focused on the sociological and the conversational necessity of borrowing and mixing (Auer 1984; Muysken 1996).

Automatic processing of mixed text: One of the first works in automatic processing of code-mixed text dates more than thirty years back (Joshi 1982), while the task of automatic language identification is even earlier (Gold 1967). Since then there have been many works in this area. For instance, (Sankoff, Poplack, and Vanniarajan 1990) reported the complexity of choosing features to determine the indicators of borrowing. This work further shows that it is not always true that only highly frequent words are borrowed, nonce words could also be borrowed along with the frequent words. According to (Field 2002), the principle of system compatibility/incompatibility can be used to determine the nature of systematic interactions of human beings. (Nzai et al. 2014) analyzed the formal conversation of Spanish-English multilingual people and found that code mixing/ word borrowing is not only restricted to daily speech but also in formal conversations. (Hadei 2016) showed that phonological integration could be evaluated to understand the phenomenon of word borrowing. In similar lines, (Sebonde 2014) showed morphological and syntactic features could be good indicators for numerical borrowings. (Senaratne 2013) reported that in many languages English words are likely to be borrowed in both formal and semi-formal text.

In parallel, works on language identification have also gained a lot of attention; for instance, (Claros, Stella, and Isharyanti 2009; Choudhury, Ahmed, and Bali 2011; Prager 1999) have analyzed the problems of language identification in code-mixed text.

Mixing in short texts and emails: There have also been works studying code-mixing in short texts like SMS. (Sotillo 2012) investigated various types of code-mixing in a corpora of 880 SMS text messages. The author observed that most often mixing takes place at the beginning of a sentence as well as through simple insertions. Similar observations about chat messages have been reported in (Bock 2013). (Negrón 2009) reported a study on code-mixing in the emails of five Spanish-English bilinguals and reiterated similar results as in the case of SMS and chat messages. However, studies of code-mixing with Chinese-English bilinguals from Hong Kong (Li 2009) and Macao (San 2009) brings forth results that contrasts the aforementioned findings and indicate that in these societies code-mixing is driven more by linguistic than social motivations.

Mixing in social media: Recently, the advent of social media has immensely propelled the research on code-mixing and word-borrowing as dynamical social phenomena. Such research has been inspired by many interesting studies on various general phenomena on language dynamics (Danescu-Niculescu-Mizil et al. 2013; Kulkarni et al. 2015).

(Hidayat 2012) noted that in Facebook, users mostly preferred inter-sentential mixing and showed that 45% of the mixing originated from real lexical needs, 40% was used for conversations on a particular topic and the rest 5% for content clarification. In contrast, (Das and Gambäck 2014) showed that in case of Facebook messages, intra-sentential mixing accounted for more than half of the cases while inter-sentential mixing accounted only for about one-third of the cases. There have also been quite a few studies on Twitter; for instance (Carter 2012) collected tweets from five different languages (Dutch, English, French, German, and Spanish), and manually inspected the multilingual micro-blogs to identify the dominant language for a specific tweet. Subsequently, a character n-gram distance metric was introduced to automate the process. In fact, in the First Workshop on *Computational Approaches to Code Switching* a shared task on code-mixing in tweets was launched and four different code-mixed corpora were collected from Twitter as a part of the shared task (Solorio et al. 2014). Language identification task has also been handled for English-Hindi and English-Bengali code-mixed tweets in (Das and Gambäck 2013). Part-of-speech tagging have been recently done for code-mixed English-Hindi tweets (Vyas et al. 2014).

Despite the presence of such a huge literature, there has not been much attempt to quantify the likeliness of borrowing of candidate foreign word in a native language. The only work that makes an attempt in this direction is (Bali et al. 2014). The biggest hurdle in doing such a study is that borrowing is a strongly social phenomenon and it is difficult to identify suitable indicators of such a lexical diffusion process unless one has access to large population level data. In this work, we show for the first time how certain simple and closely related signals encoding the language usage of social meida users can help us construe appropriate metrics to quantify the likeliness of borrowing of a foreign word. We show, through a series of rigorous experiments how our metrics by far beats the most competitive baseline reported in the literature.

Datasets and preprocessing

We consider English-Hindi code-mixed tweets from India for the purpose of our experiments. To bootstrap the data collection process, we crawl tweets (between Nov 2015 and Jan 2016) related to 28 hashtags representing different Indian contexts covering important topics such as sports, religion, movies, politics etc. This process results in the collection of 811981 tweets. We language-tag (see details later in this section) each tweet so crawled and find that there are 3982 users who use mixed language for tweeting. We then systematically crawl the time lines of these 3982 users between Feb 2016 and March 2016 to gather more mixed language tweets. Using this two step process we collect a total of 1550714 distinct tweets. From this data, we filter out tweets that are not written in romanized script, tweets having only URLs and tweets having empty content. Post filtering we obtain 787606 tweets which we use for the rest of the analysis.

Language tagging: We tag each word in a tweet with the language of its origin using the method outlined in (Gella,

Sharma, and Bali 2013). The different tags that a word can have are: *En* (English), *Hi* (Hindi), *NE* (Named Entity) and *Others* (hash-tags, URLs, twitter user names etc.). Based on the word level tag, we create a tweet level tag as follows:

1. *En*: Almost every word ($> 90\%$) in the tweet is tagged as *En*.
2. *Hi*: Almost every word ($> 90\%$) in the tweet is tagged as *Hi*.
3. *CME*: Code-mixed tweet but majority (i.e., $> 50\%$) of the words are tagged as *En*.
4. *CMH*: Code-mixed tweet but majority (i.e., $> 50\%$) of the words are tagged as *Hi*.
5. *CMEQ*: Code-mixed tweet having equal number of words tagged as *En* and *Hi* respectively.
6. *Code Switched*: There is a trail of Hindi words followed by a trail of English words or vice versa.

In table 1 we note the number and percentage of tweets in each of the above six categories in which the tweets are labeled. Like the word level, the tagger also provides a phrase level language tag. Once again, the different tags that an entire phrase can have are: *En*, *Hi* and *Oth* (Other). We shall use these word and phrase level tags in order to define our metrics in the next section.

Frequent foreign words: In this step we compute the most frequent foreign (i.e., English words) in our tweet corpus. Since we are interested in the frequency of the English word only when it appears as a foreign word we do not consider the (i) *Hi* tweets since they do not have any foreign word, (ii) *En* tweets since here the English words are not foreign words and the (iii) code-switched tweets. Based on the frequency of usage of English as a foreign word, we select the top 1000 English words. Removal of stop words and text normalization leaves beyond 230 nouns. We note these words in the box below:

'welfare', 'anniversary', 'tribute', 'box', 'victory', 'thing', 'lot', 'youth', 'need', 'nation', 'birth', 'people', 'muslims', 'god', 'water', 'teacher', ''airport', 'army', 'room', 'answer', 'blood', 'law', 'light', 'chief', 'green', 'office', 'border', 'food', 'university', 'side', 'event', 'health', 'reason', 'city', 'station', 'theatre', 'crore', 'ground', 'college', 'bomb', 'corruption', 'court', 'opposition', 'respect', 'life', 'air', 'rail', 'student', 'government', 'mom', 'aunty', 'weekend', 'age', 'protest', 'guy', 'company', 'bollywood', 'place', 'message', 'friend', 'mind', 'mobile', 'view', 'volunteer', 'moment', 'rest', 'suicide', 'lyrics', 'group', 'death', 'home', 'way', 'brother', 'house', 'blue', 'wedding', 'reaction', 'terrorist', 'person', 'mother', 'press', 'election', 'power', 'question', 'lord', 'birthday', 'president', 'half', 'day', 'internet', 'number', 'service', 'morning', 'waste', 'voice', 'evening', 'night', 'luck', 'son', 'favourite', 'captain', 'video', 'sun', 'body', 'experience', 'family', 'use', 'music', 'date', 'teaser', 'share', 'man', 'paper', 'lunch', 'logo', 'season', 'job', 'game', 'post', 'gift', 'poster', 'film', 'test', 'performance', 'price', 'plan', 'class', 'shot', 'report', 'prime', 'state', 'exam', 'success', 'road', 'form', 'problem', 'check', 'wife', 'boy', 'car', 'heart', 'scam', 'style', 'police', 'issue', 'card', 'country', 'boss', 'party', 'entry', 'uncle', 'politics', 'father', 'parliament', 'work', 'sunday', 'story', 'play', 'request', 'week', 'playlist', 'matter', 'superstar', 'traffic', 'suit', 'woman', 'cool', 'history', 'money', 'bat', 'seat', 'score', 'photo', 'parents', 'decision', 'girlfriend', 'picture', 'month', 'song', 'word', 'school', 'hero', 'degree', 'love', 'train', 'end', 'wrong', 'main', 'scene', 'bank', 'miss', 'king', 'channel', 'face', 'link', 'news', 'media', 'mood', 'book', 'selfie', 'bus', 'status', 'petrol', 'railway', 'budget', 'well', 'development', 'team', 'phone', 'baby', 'sir', 'interview', 'fan', 'trailer', 'year', 'girl', 'time', 'review', 'madam', 'movie', 'minister', 'joke', 'century', 'cup', 'match', 'world', 'temple', 'wicket', 'cricket', 'star'

Table 1: Number and percentage of tweets in each of the six categories in which the tweets are labeled.

Type	Number of Tweets	percentage
En	645655	81.97
Hi	24960	3.16
CME	31998	4.06
CMH	39877	5.06
CMEQ	3584	0.455
CS	41532	5.27

Newspaper dataset: As we shall see, for the construction of the baseline ranking we shall need to resort to counting the frequency of the foreign words (i.e., English words) and their Hindi translations in a newspaper corpus as has been outlined in (Bali et al. 2014). For this purpose, we use the FIRE dataset built from the Hindi Jagaran newspaper corpus[1] which is written in Devanagari script.

Methodology

In this section, we present the baseline metric and our metrics based on the population level signals to quantify the likeliness of borrowing. We also present a scheme for sampling an appropriate set of target words for the purpose of our experiments.

Baseline metric and ranking

Baseline metric – We consider the $log(\frac{F_E}{F_H})$ value proposed in (Bali et al. 2014) as the baseline metric. Here F_E denotes the frequency of the Devanagari transliterated form of the word w in the Jagaran corpus. F_H, on the other hand, denotes the frequency of the Hindi translation of the word w in the Jagaran corpus. Both the transliteration and the translation of the words have been done by a set of volunteers who are native Hindi speakers. The authors in (Bali et al. 2014) claim that the more positive the value of this metric is for a word w, the higher is the likeliness of its being borrowed. The more negative the value is, the higher are the chances that the word w is an instance of code-mixing.

Ranking – Based on the values obtained from the above metric for a set of target words, we rank these words; words with high positive values feature at the top of the rank list and words with high negative values feature at the bottom of the list. For two words having the same $log(\frac{F_E}{F_H})$ value, we resolve the conflict by assigning each of these the average of their two rank positions. In the next section, we shall compare this rank list with the one obtained from the ground truth responses.

Proposed metric and ranking

In this section, we present three novel and closely related metrics based on the language usage patterns of the users of social media.

[1]Jagaran corpus: **http:/fire.irsi.res.in/fire/static/data**

Unique User Ratio (UUR) – The Unique User Ratio for word usage across languages is defined as follows:

$$UUR(w) = \frac{U_{Hi} + U_{CMH}}{U_{En}} \quad (1)$$

where U_{Hi} is the number of unique users who have used the word w in a Hindi tweet at least once, U_{En} is the number of unique users who have used the word w in an English tweet at least once and U_{CMH} is the number of users who have used the word w in a code-mixed Hindi tweet at least once. Higher the value of UUR higher should be the likeliness of the word w being borrowed.

Unique Tweet Ratio (UTR) – The Unique Tweet Ratio for word usage across languages is defined as follows:

$$UTR(w) = \frac{T_{Hi} + T_{CMH}}{T_{En}} \quad (2)$$

where T_{Hi} is the total number of Hindi tweets which contain the word w, T_{En} is the total number of English tweets which contain the word w and T_{CMH} is the total number of CMH tweets which contain the word w . Higher the value of UTR higher should be the likeliness of the word w being borrowed.

Unique Phrase Ratio (UPR) – The Unique Phrase Ratio for word usage across languages is defined as follows:

$$UPR(w) = \frac{P_{Hi}}{P_{En}} \quad (3)$$

where P_{Hi} is the number of Hindi phrases which contain the word w, P_{En} is the number of English phrases which contain the word w. Note that unlike the definitions of UUR and UTR that exploit the word level language tags, the definition of UPR exploits the phrase level language tags. Once again, higher the value of UPR higher should be the likeliness of the word w being borrowed.

Ranking – We prepare a separate rank list of the target words based on each of the three proposed metrics – UUR, UTR and UPR. In the next section, we shall compare these rank lists with the one prepared from the ground truth responses.

Target word selection

We use following method to select the final set of *target words* from the 230 nouns for our evaluation.

Clustering In language processing, context plays an important role in understanding different properties of a word. For our study, we also attempt to use the language tags as features of the context words for a given *target word*. Our hypothesis here is that there should exist classes of words that have similar context features and the likelihood of being borrowed in each class should be different. For example, when an English word is surrounded by mostly Hindi words it seems to be more likely borrowed. We present two examples in the box below to illustrate this.

Example I:
@KapilianPooja Welcome. ***Film*** *jaroor dekhna. Nahi to injection ready hai.*
- Dr Mashoor Gulati (@ItsDrGulati)
Example II:
@HelpU_Trust @DrKumarVishwas
Kuch to ache se karo sirji....
Har jagah bhaagte rehna is not a good ***thing***.
- lovely sethi(@lovelysethii)

In *Example I* the English word "film" is surrounded by mostly Hindi words. On the other hand, in *Example II* the English word "thing" is surrounded mostly by English words. Note that the word "film" is very commonly used by Hindi monolingual speakers and is therefore highly likely to have been borrowed unlike the English word "thing" which is arguably an instance of mixing. This socio-linguistic difference seems to be very appropriately captured by the language tag of the surrounding words of these two words in the respective tweets.

Therefore, as a first step toward sampling a set of representative *target words*, we cluster the list of 230 words into contextually similar groups as follows.

Construction of feature vectors – We represent a context feature for a *target word* as a tuple $\{P_b, P_a\}$ where P_b is the language tag for the word before the *target word* (i.e., the left context) and P_a is the language tag of the word after the *target word* (i.e., the right context). Each of P_b and P_a can be either "E" indicating English, "H" indicating Hindi or "$" indicating the boundary (i.e., beginning or end) of the tweet. Thus, we have *eight* feature combinations of the left and the right contexts of a *target word* – "EE", "HH", "EH", "HE", "$E", "E$", "$H", "H$" while "$$" is not possible. For every *target word*, we compute the percentage of occurrences of each of these combinations.

Note that we compute these percentages from the three different categories of tweets – *CME*, *CMH* and *CMEQ*. Thus, for every *target word* we have a final feature vector of length 24, each entry denoting the percentage of one feature combination in a particular tweet category. We show example feature vectors for some words in figure 2.

K-means clustering – We use the feature representation of the words to cluster them into contextually similar groups. We use K-means clustering (Hartigan and Wong 1979) for the purpose of clustering. Hence we vary the value of K and using the traditional elbow method (Tibshirani, Walther, and Hastie 2001) we obtain 15 as the optimal value of K. This process therefore groups the 230 nouns into 15 different clusters.

Final selection of target words In this final stage, we select two sets of target words using two different strategies as follows.

Baseline biased words (bbw) – We compute $log(\frac{F_E}{F_H})$ values for all the words in each cluster. From each cluster, we select two words having the highest and the lowest values of $log(\frac{F_E}{F_H})$. This constitutes a set of 30 words. Note that this selection is biased toward the baseline metric discussed

Figure 2: Stacked plot representing feature vectors of four different words. Note that the feature vectors of the word pairs (i) "job" and "film" and (ii) "moment" and "protest" are very similar. The fractional counts of the eight combinations for each tweet category should sum up to one; since there are three tweet categories so the total size of the stacked plot is three.

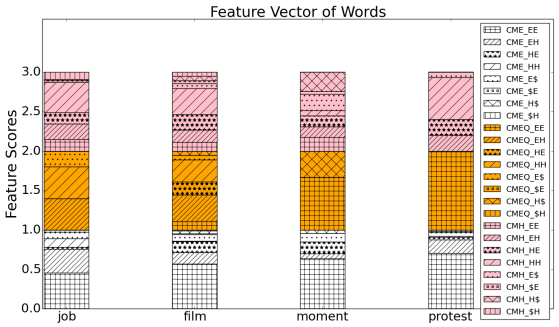

above so as to give maximum possible advantage to the baseline ranking.

Randomly selected words (ran) – In order to check how well all the proposed metrics behave in general, i.e., to investigate the recall of each metric, we randomly select 27 words from the 15 clusters. These 27 words are kept completely different from the baseline biased words.

Full set of words ($full$) – Thus, in total we select 57 target words for the purpose of our evaluation. We present these words in the box below.

Baseline biased words – 'thing', 'way', 'woman', 'press', 'wrong', 'well', 'matter', 'reason', 'question', 'guy', 'moment', 'week', 'luck', 'president', 'body', 'job', 'car', 'god', 'gift', 'status', 'university', 'lyrics', 'road', 'politics', 'parliament', 'review', 'scene', 'seat', 'film', 'degree'
Randomly selected words – 'people', 'play', 'house', 'service', 'rest', 'boy', 'month', 'money', 'cool', 'development', 'group', 'friend', 'day', 'performance', 'school', 'blue', 'room', 'interview', 'share', 'request', 'traffic', 'college', 'star', 'class', 'superstar', 'petrol', 'uncle'

Results and Discussion

Evaluation criteria

In this section, we present extensive evaluation to demonstrate the effectiveness of the three proposed metrics UUR, UTR and UPR in quantifying the likeliness of borrowing. In particular, we present a four step approach as follows. We measure (i) how well the UUR, UTR and UPR based ranking of the bbw set, the ran set and the $full$ set correlate with the ground truth ranking (discussed in the next section) in comparison to the rank given by the baseline metric, (ii) how well the different rank ranges obtained from our metric align with the ground truth as compared to the baseline metric, (iii) whether there are some systematic effects of the age group of the survey participants on the rank correspondence, and (iv) how our metrics if computed from the tweets of users who (a) rarely mix languages, (b) almost always mix languages and (c) are in between (a) and (b), align with the ground truth.

Rank correlation: We measure the standard Spearman's rank correlation (ρ) (Zar 1972) pairwise between rank lists generated by (i) UUR and ground truth, (ii) UTR and ground truth, (iii) UPR and ground truth, and (iv) baseline and ground truth.

We shall describe the next three measurements taking UUR as the running example. The same can be extended verbatim for the other two similar metrics.

Rank ranges: We split each of the three rank lists (UUR, ground truth and baseline) into five different equal-sized ranges as follows – (i) surely borrowed (SB) containing top 20% words from each list, (ii) likely borrowed (LB) containing the next 20% words from each list, (iii) borderline (BL) constituting the subsequent 20% words from each list, (iv) likely mixed (LM) comprising the next 20% words from each list and (v) surely mixed (SM) having the last 20% words from each rank list. Therefore, we have three sets of five buckets, one set each for UUR, the ground truth and the baseline based rank list.

Next we calculate the bucket-wise correspondence between (i) the UUR and the ground truth set and (ii) the baseline and the ground truth set in terms of standard *precision* and *recall* measures. For our purpose, we adapt these measures as follows.

G: ground truth bucket set, B_b: baseline bucket set, U_b: UUR bucket set;
$BS \in \{B_b, U_b\}$, T (type of bucket) = {SB, LB, BL, LM, SM};
b_t = words in type t bucket from BS, g_t = words in type t bucket from G, $t \in T$;
tp_t (no. of true positives) = $|b_t \cap g_t|$, fp_t (no. of false positives) = $|b_t - g_t|$, tn_t (no. of true negatives) = $|g_t - b_t|$;
Bucket-wise *precision* and *recall* are then given by:
$precision(b_t) = \frac{tp_t}{fp_t + tp_t}$
$recall(b_t) = \frac{tp_t}{tn_t + tp_t}$
For a given set, we obtain the overall *macro precision* (*recall*) by averaging the *precision* (*recall*) values over the five buckets. For a given set, we also obtain the overall *micro precision* by first adding the true positives across all the buckets and then normalizing by the sum of the true and the false positives over all the buckets. We take an equivalent approach for obtaining the *micro recall*.

Age group effect: Here we construct two ground truth rank lists one using the responses of the participants with age below 30 (young population) and the other using the responses of the rest of the participants (elderly population). Next we repeat the above two evaluations considering each of the new ground truth rank lists.

Extent of language mixing: Here we divide all the 3982 users into three categories – (i) High (users who have more than 20% of tweets as code-mixed), (ii) Mid (users who have 7–20% of their tweets as code-mixed, and (iii) Low (users who have less than 7% of their tweets as code-mixed). We create three UUR based rank lists for each of these three user categories and respectively compare them with the ground truth rank list.

Ground truth preparation

Since it is very difficult to obtain a suitable ground truth to validate the effectiveness of our proposed ranking scheme, we launched an online survey to collect human judgment for each of the 57 *target words*.

Online survey We conducted the online survey[2] among 58 volunteers majority of whom were either native Hindi speakers or had very high proficiency in reading and writing Hindi. The participants were selected from different age groups and different educational backgrounds. Every participant was asked to respond to a multiple choice question about each of the 57 *target words*. Therefore, for every single *target word*, 58 responses were gathered. The multiple choice question had the following three options and the participants were asked to select the one they preferred the most and found more natural – (i) a Hindi sentence with the *target word* as the only English word, (ii) the same Hindi sentence in (i) but with the *target word* replaced by its Hindi translation and (iii) none of the above two options. There were no time restrictions imposed while gathering the responses, i.e., the volunteers theoretically had unlimited time to decide their responses for each *target word*.

Language preference factor For each *target word*, we compute a *language preference factor* (LPF) defined as $(Count_{En} - Count_{Hi})$, where $Count_{Hi}$ refers to the number of survey participants who preferred the sentence containing the Hindi translation of the *target word* while $Count_{En}$ refers to the number of survey participants who preferred the sentence containing the *target word* itself. More positive values of LPF denotes higher usage of *target word* as compared to its Hindi translation and therefore higher likeliness of the word being borrowed.

Ground truth rank list generation We generate the ground truth rank list based on the LPF score of a *target word*. The word with the highest value of LPF appears at the top of the ground truth rank list and so on in that order. Tie breaking between *target words* having equal LPF values is done by assigning average rank to each of these words.

Age group based rank list: As discussed in the previous section, we prepare the age group based rank lists by first splitting the responses of the survey participants in two groups based on their age – (i) young population (age < 30) and (ii) elderly population (age ≥ 30). For each group we then construct a separate LPF based ranking of the *target words*.

Correlation among rank lists

The Spearman's rank correlation coefficient (ρ) of the rank lists for the bbw set, the ran set and the $full$ set according to the baseline metric, UUR, UTR and UPR with respect to the ground truth metric LPF are noted in 2. We observe that for the $full$ set, the ρ between the rank lists obtained from all the three metrics UUR, UTR, and UPR with respect to the ground truth is ~ 0.62 which is more than double the ρ (~ 0.26) between the baseline and the ground truth rank list. This clearly shows that the proposed metrics are able to identify the likeliness of borrowing quite accurately and far better than the baseline. Further, a remarkable observation is that our metrics outperform the baseline metric even for the bbw set that is baseline-biased. Likewise, for the ran set, our metrics outperform the baseline indicating a superior recall on arbitrary words.

In order to help the readers to better visualize, we show in figure 3 the ranking of the $full$ set of words obtained from the ground truth, the baseline and the UUR metric. The figure clearly shows that UUR correlates with the ground truth far better than the baseline. Very similar results are obtained for the two other mertics UTR and UPR and therefore not shown.

We present the subsequent results for the $full$ set and the UUR metric. The results obtained for the other two metrics UTR and UPR are very similar and therefore not shown.

Table 2: Spearman's rank correlation coefficient (ρ) among the different rank lists. Best result is marked in bold.

Rank-List$_1$	Rank-List$_2$	$\rho - bbw$	$\rho - ran$	$\rho - full$
UUR	Ground truth	**0.67**	**0.64**	**0.62**
UTR	Ground truth	**0.66**	**0.63**	**0.63**
UPR	Ground truth	**0.66**	**0.64**	**0.62**
Baseline	Ground truth	0.49	0.14	0.26

Table 3: Number of words falling in each bucket of three bucket sets.

Bucket type	Ground truth bucket	Baseline bucket	UUR bucket
SB	11	11	11
LB	11	11	11
BL	12	12	12
LM	11	11	11
SM	12	12	12

Rank list alignment across rank ranges

The number of *target words* falling in each bucket across the three rank lists are same and are noted in table 3. Thus, the *precision* and *recall* as per the definition are also the same. The bucket-wise *precision*/*recall* for the baseline and UUR with respect to the ground truth are noted in table 4. We observe that while in the SB bucket both the baseline and UUR perform equally well, for all the other buckets UUR massively outperforms the baseline. This implies that for the case where the likeliness of borrowing is the strongest, the baseline does as good as UUR. However, as one moves down the rank list, UUR turns out to be a considerably better predictor than the baseline. The overall *macro* and *micro precision*/*recall* as shown in table 5 further strengthens our observation that UUR is a better metric than the baseline.

Age group based analysis

As already discussed earlier, we split the ground truth responses based on the age group of the survey participants. In particular, we wish to observe the effect of this split on our

[2]Survey portal: **https://goo.gl/forms/L0kJm8BNMhRj0jA53**

Figure 3: The rank ordered histograms of *target words* ranked by various metrics.

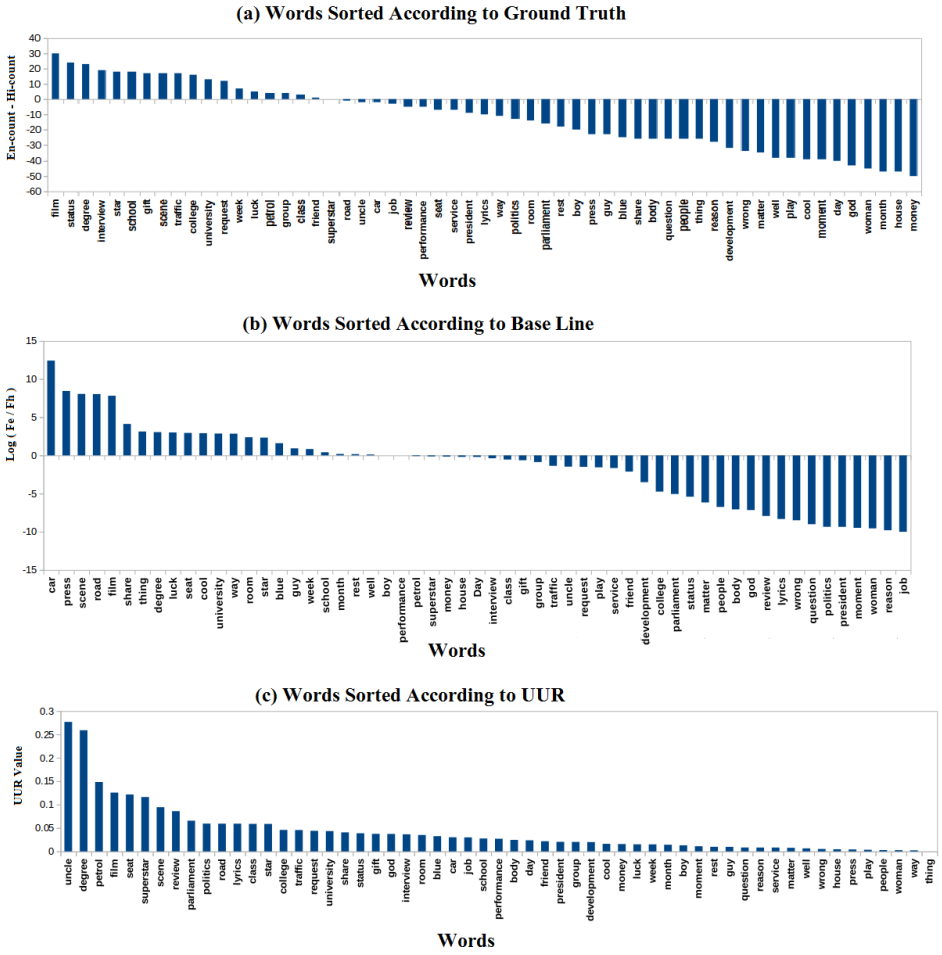

Table 4: Bucket-wise *precision/recall*. Best results are marked in bold.

Bucket type	*prec./rec.* Baseline	*prec./rec.* UUR
SB	**0.27**	**0.27**
LB	0.09	**0.18**
BL	0.08	**0.33**
LM	0.18	**0.36**
SM	0.33	**0.50**

Table 5: Overall *macro* and *micro precision/recall*. Best results are marked in bold.

Measure	Baseline	UUR
Macro prec./rec.	0.19	**0.33**
Micro prec./rec.	0.19	**0.33**

Figure 4: Age distribution of the survey participants.

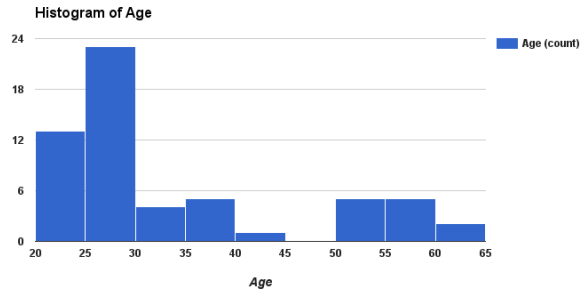

ranking results. The distribution of participants of different age groups is shown in figure 4. Based on this distribution we split the responses into two groups as discussed earlier – (i) young population (age < 30) and (ii) elderly population (age ≥ 30).

Rank correlation: The Spearman's rank correlation of

Table 6: Spearman's rank correlation across the two age groups. Best results are marked in bold.

Rank-List$_1$	Rank-List$_2$	ρ
Baseline	Ground-truth-Young	0.26
UUR	Ground-truth-Young	**0.62**
Baseline	Ground-truth-Elder	0.27
UUR	Ground-truth-Elder	**0.53**

Table 7: Bucket-wise *precision/recall* for UUR and the $baseline$ metrics for the two new ground truths. Best results are marked in bold.

Bucket type	Young-Baseline *prec./Rec.*	Young-UUR *prec./rec*	Elder-Baseline *prec.*	Elder-UUR *prec.*	Elder-baseline *rec.*	Elder-UUR *rec.*
SB	**0.27**	**0.27**	0.27	**0.36**	0.25	**0.33**
LB	0.09	**0.18**	0.09	**0.18**	0.08	**0.17**
BL	0.08	**0.33**	**0.16**	0.08	**0.28**	0.14
LM	0.18	**0.36**	0.18	**0.45**	0.14	**0.35**
SM	0.33	**0.5**	**0.41**	0.25	**0.41**	0.25

UUR and the baseline rank lists with these two ground truth rank lists are shown in table 6. Interestingly, the correlation between UUR rank list and the young population ground truth is better than the elderly population ground truth. This possibly indicates that UUR is able to predict recent borrowings more accurately. However, note that the UUR rank list has a much higher correlation with both the ground truth rank lists as compared to the baseline rank list.

Rank ranges: Table 7 shows the bucket-wise *precision* and *recall* for UUR and the baseline metrics with respect to two new ground truths. For the young population once again the number of words in each bucket for all the three sets is the same thus making the values of the *precision* and the *recall* same. In fact, the *precision/recall* for this ground truth is exactly same as in the case of the original ground truth.

In contrast, when we consider the ground truth based on the responses of the elderly population, the number of words across the different buckets are different across the three sets. In this case, we observe that the *precision/recall* values are better for the UUR metric in SB, LB and LM buckets.

Finally, the overall *macro* and *micro precision* and *recall* for both the age groups are noted in table 8. Once again, for both the young and the elderly population based ground truths, the *macro* and *micro precision* and *recall* values for the UUR metric are higher compared to that of the baseline.

Table 8: Overall *macro* and *micro precision* and *recall* for the two new ground truths. Best results are marked in bold.

Measure	Young-Baseline	Young-UUR	Elder-Baseline	Elder-UUR
Macro prec.	0.19	**0.33**	0.22	**0.27**
Macro rec.	0.19	**0.33**	0.23	**0.25**
Micro prec./rec.	0.19	**0.33**	0.23	**0.26**

Table 9: Spearman's correlation between UUR and the ground truth in the different user buckets. Best results are marked in bold.

Bucket	Number of users	ρ
High	402	0.52
Mid	844	0.60
Low	2736	**0.65**

Results based on the extent of language mixing

As mentioned earlier, we divide the set of 3982 users into three categories. The Spearman's correlation between UUR and the ground truth for each of these buckets are given in table 9. As we can see, for Low bucket the ρ value is maximum. This points to the fact that the signals of borrowing is strongest from the users who rarely mix languages.

Conclusion

In this paper we presented exhaustive experiments to quantify the likeliness of borrowing of a set of appropriately sampled *target words*. Some of our key contributions are

- We introduced a context based classification method to appropriately sample a list of 57 English *target words*.
- We launched an extensive online survey among 58 human judges to prepare a ground-truth data indicating the likeliness of borrowing of a *target word*.
- As a central contribution, we defined three closely related novel metrics based on social media signals indicating the likeliness that a *target word* shall be borrowed.

Some of our key results from this study are:

- Remarkably, the rankings based on our metrics are more than two times more correlated to the ground-truth ranking compared to that of the baseline.
- Further, our metrics predict the likeliness of borrowing better for the young population (age group below 30) which possibly indicates that these are able to capture the early signals of borrowing.
- Finally, if we compute the metrics from the tweets of those users who seldom code-mix, we obtain the best signals of borrowing.

There are quite a few interesting directions of future research. First, we would like to obtain more theoretical insights into the better performance of these metrics by possibly having a dynamical model of word borrowing. Further, we would also like to incorporate our findings into standard tasks of multilingual IR and NLP and test if this leads to systematic performance enhancements.

References

[Auer 1984] Auer, P. 1984. *Bilingual Conversation*. John Benjamins.

[Bali et al. 2014] Bali, K.; Sharma, J.; Choudhury, M.; and Vyas, Y. 2014. "i am borrowing ya mixing?" an analysis of english-hindi code mixing in facebook. In *First workshop on Computational approaches to code-switching, EMNLP*, 116.

[Barman et al. 2014] Barman, U.; Das, A.; Wagner, J.; and Foster, J. 2014. Code mixing: A challenge for language identification in the language of social media. *EMNLP 2014* 13.

[Bock 2013] Bock, Z. 2013. Cyber socialising: Emerging genres and registers of intimacy among young south african students. *Language Matters: Studies in the Languages of Africa* 44(2):68–91.

[Carter, Weerkamp, and Tsagkias 2013] Carter, S.; Weerkamp, W.; and Tsagkias, M. 2013. Micro-blog language identification: Overcoming the limitations of short, unedited and idiomatic text. *Language Resources and Evaluation* 47(1):195–215.

[Carter 2012] Carter, S. 2012. Exploration and exploitation of multilingual data for statistical machine translation. In *PhD thesis*. University of Amsterdam.

[Choudhury, Ahmed, and Bali 2011] Choudhury, M.; Ahmed, U. Z.; and Bali, K. 2011. Challenges in designing input method editors for indian languages: The role of word-origin and context. *Advances in Text Input Methods (WTIM 2011)* 1.

[Claros, Stella, and Isharyanti 2009] Claros, C.; Stella, M.; and Isharyanti, N. 2009. Code-switching and code-mixing in internet chatting: Between 'yes', 'ya', and 'si'- a case study. *The Jalt Call Journal* 5(3):67–78.

[Danescu-Niculescu-Mizil et al. 2013] Danescu-Niculescu-Mizil, C.; West, R.; Jurafsky, D.; Leskovec, J.; and Potts, C. 2013. No country for old members: User life-cycle and linguistic change in online communities. In *WWW*, 307–318.

[Das and Gambäck 2013] Das, A., and Gambäck, B. 2013. code-mixing in social media text: The last language identification frontier? *TAL* 54(3):41–64.

[Das and Gambäck 2014] Das, A., and Gambäck, B. 2014. Identifying languages at the word level in code-mixed indian social media text. In *ICON*, 169–178.

[Field 2002] Field, F. W. 2002. *Linguistic borrowing in bilingual contexts*, volume 62. John Benjamins Publishing.

[Gella, Sharma, and Bali 2013] Gella, S.; Sharma, J.; and Bali, K. 2013. Query word labeling and back transliteration for indian languages: Shared task system description. *FIRE Working Notes* 3.

[Gold 1967] Gold, E. M. 1967. Language identification in the limit. *Information and Control* 10(5):447–474.

[Hadei 2016] Hadei, M. 2016. Single word insertions as code-switching or established borrowing? *International Journal of Linguistics* 8(1):14–25.

[Hartigan and Wong 1979] Hartigan, J. A., and Wong, M. A. 1979. Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28(1):100–108.

[Hidayat 2012] Hidayat, T. 2012. An analysis of code switching used by facebookers (a case study in a social network site). In *BA thesis*. English Education Study Program, College of Teaching and Education (STKIP), Bandung, Indonesia.

[Joshi 1982] Joshi, A. K. 1982. Processing of sentences with intra-sentential code-switching. In *Coling*, 145–150.

[Kulkarni et al. 2015] Kulkarni, V.; Al-Rfou, R.; Perozzi, B.; and Skiena, S. 2015. Statistically significant detection of linguistic change. In *WWW*, 625–635.

[Li 2009] Li, D. C. S. 2009. Cantonese-english code-switching research in hong-kong: a y2k review. *World Englishes* 19(3):305–322.

[Muysken 1996] Muysken, P. 1996. Code-switching and grammatical theory. In Milroy, L., and Muysken, P., eds., *One speaker, two languages: Cross-disciplinary perspectives on code-switching*, 177–198. Cambridge University Press.

[Myers-Scotton 2002] Myers-Scotton, C. 2002. *Contact linguistics: Bilingual encounters and grammatical outcomes*. Oxford University Press.

[Negrón 2009] Negrón, R. G. 2009. Spanish-english code-switching in email communication. *Language@Internet* 6(3).

[Nzai et al. 2014] Nzai, V. E.; Feng, Y.-L.; Medina-Jiménez, M. R.; and Ekiaka-Oblazamengo, J. 2014. Understanding code-switching & word borrowing from a pluralistic approach of multilingualism.

[Pérez-Sabater 2012] Pérez-Sabater, C. 2012. The linguistics of social networking: A study of writing conventions on facebook. *Linguistik Online* 56(6).

[Prager 1999] Prager, J. M. 1999. Linguini: Language identification for multilingual documents. *Journal of Management Information Systems* 16(3):71–101.

[Rudra et al. 2016] Rudra, K.; Rijhwani, S.; Begum, R.; Bali, K.; Choudhury, M.; and Ganguly, N. 2016. Understanding language preference for expression of opinion and sentiment: What do hindi-english speakers do on twitter? In *EMNLP*.

[San 2009] San, H. K. 2009. Chinese-english code-switching in blogs by macao young people. In *MSc. thesis*. University of Edinburgh.

[Sankoff, Poplack, and Vanniarajan 1990] Sankoff, D.; Poplack, S.; and Vanniarajan, S. 1990. The case of the nonce loan in tamil. *Language variation and change* 2(01):71–101.

[Sebonde 2014] Sebonde, R. Y. 2014. Code-switching or lexical borrowing: Numerals in chasu language of rural tanzania. *Journal of Arts and Humanities* 3(3):67.

[Senaratne 2013] Senaratne, C. D. W. 2013. Borrowings or code mixes: The presence of lone english nouns in mixed discourse.

[Solorio et al. 2014] Solorio, T.; Blair, E.; Maharjan, S.; Bethard, S.; Diab, M.; Gohneim, M.; Hawwari, A.; AlGhamdi, F.; Hirschberg, J.; Chang, A.; and Fung, P. 2014. Overview for the first shared task on language identification in code-switched data. In *EMNLP First Workshop on Computational Approaches to Code Switching*, 62–72.

[Sotillo 2012] Sotillo, S. 2012. Ehhhh utede hacen plane sin mi???:@ im feeling left out :(form, function and type of code switching in sms texting. In *ICAME*, 309–310.

[Thomason 2003] Thomason, S. G. 2003. Contact as a source of language change. *The handbook of historical linguistics* 687–712.

[Tibshirani, Walther, and Hastie 2001] Tibshirani, R.; Walther, G.; and Hastie, T. 2001. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(2):411–423.

[Vyas et al. 2014] Vyas, Y.; Gella, S.; Sharma, J.; Bali, K.; and Choudhury, M. 2014. POS tagging of english-hindi code-mixed social media content. In *EMNLP*, volume 14, 974–979.

[Zar 1972] Zar, J. H. 1972. Significance testing of the spearman's rank correlation coefficient. *Journal of the American Statistical Association* 67(339):578–580.